A Cost-Benefit Analysis of a Campus Computing Grid

PRESTON SMITH
ROSEN CENTER FOR ADVANCED COMPUTING
PURDUE UNIVERSITY
PSMITH@PURDUE.EDU

Condor Week 2011

Overview

- Introduction
  - The Problem
  - Significance of the Problem
- Methodology
  - Costs
    - Benchmarking
  - Capacity
  - Utility
- Findings
- Conclusions

Background

- Purdue University Campus Grid
  - Large, high throughput, computation resource – 42,000 processor cores
- Frequently linked to efforts to reduce IT costs
  - Claims include
    - Power savings, maximizing investment in IT
    - HPC resource using existing equipment
    - *No marginal cost increase*

The Problem

- What is the Additional Cost of Having a Campus Grid?
  - On top of existing IT investment
  - People say it's basically zero – but how close is it in reality?
- An institution needs information for designing an HPC resource
  - Therefore, I define a model for identifying the costs and benefits of building a campus grid

Significance of the Problem

Significance of the Problem

- Appropriate Computations
  - IU study reports:
    - 66% of all jobs on TeraGrid in 2004-2006 were single-CPU jobs
    - 80% of those jobs ran for two hours or less
  - Purdue University
    - 35.4 million single-core serial jobs in 2008-10
    - Average runtime of 1.35 hours
      - **This is 21% of all HPC hours consumed at Purdue.**

A large amount of work is appropriate for a campus grid

Significance of the Problem

- Size of Grid Resource
  - 27,000 desktop machines at Purdue
    - 2 cores per machine – 54,000 cores on desktops
  - 30,000 cores of HPC clusters
  - **84,000 cores potentially usable by the grid**
    - **40,000 used by the grid today**

- Only 17 systems on 2010 Top 500 with more than 40,000 cores!
  - 200 TF theoretical performance – top 20 machine

Significance of the Problem

- Power Cost of Desktop Computers
  - 111W idle
  - 160W at full load
- Purdue's 27,000 desktops
  - 2.99 MW/hour, for a total of 26,253.7 MW per year
- Idle to fully loaded
  - Estimated additional cost of $393,805.80 per year

Methodology

- Identify and calculate baseline costs
  - Clusters
  - Desktop, student lab IT
- Identify and calculate additional costs
  - Staff, power, hardware
- Measure capacity of the grid
  - Sample the state of the grid over 2-week period
- Benchmark
  - Condor nodes
  - Amazon EC2
- Normalize Costs
  - To Amazon EC2
- Collect and Report Output of Grid
  - Cost per productivity metric

Pre-Normalized Costs

	Per Hour Cost
Labs	$0.0445
Steele	$0.0218
Coates	$0.0237
Condor	$0.03
EC2	$0.17

- Labs, Steele, and Coates are all derived from Purdue TCO data
- "Condor" is average of all three
- "EC2" is retail price (per core) of EC2 "Large" instance

Normalization

- Internal benchmarks
  - Condor runs and presents two predefined benchmarks
    - Kflops (LINPACK)
    - MIPS (Dhrystone)
      - Are these benchmarks meaningful enough to normalize cost?
- Application benchmark
  - Use a benchmark that relates to real performance of an application
    - NAS Parallel Benchmarks
    - Single CPU BT, SP, Class C

Normalization - Benchmarks

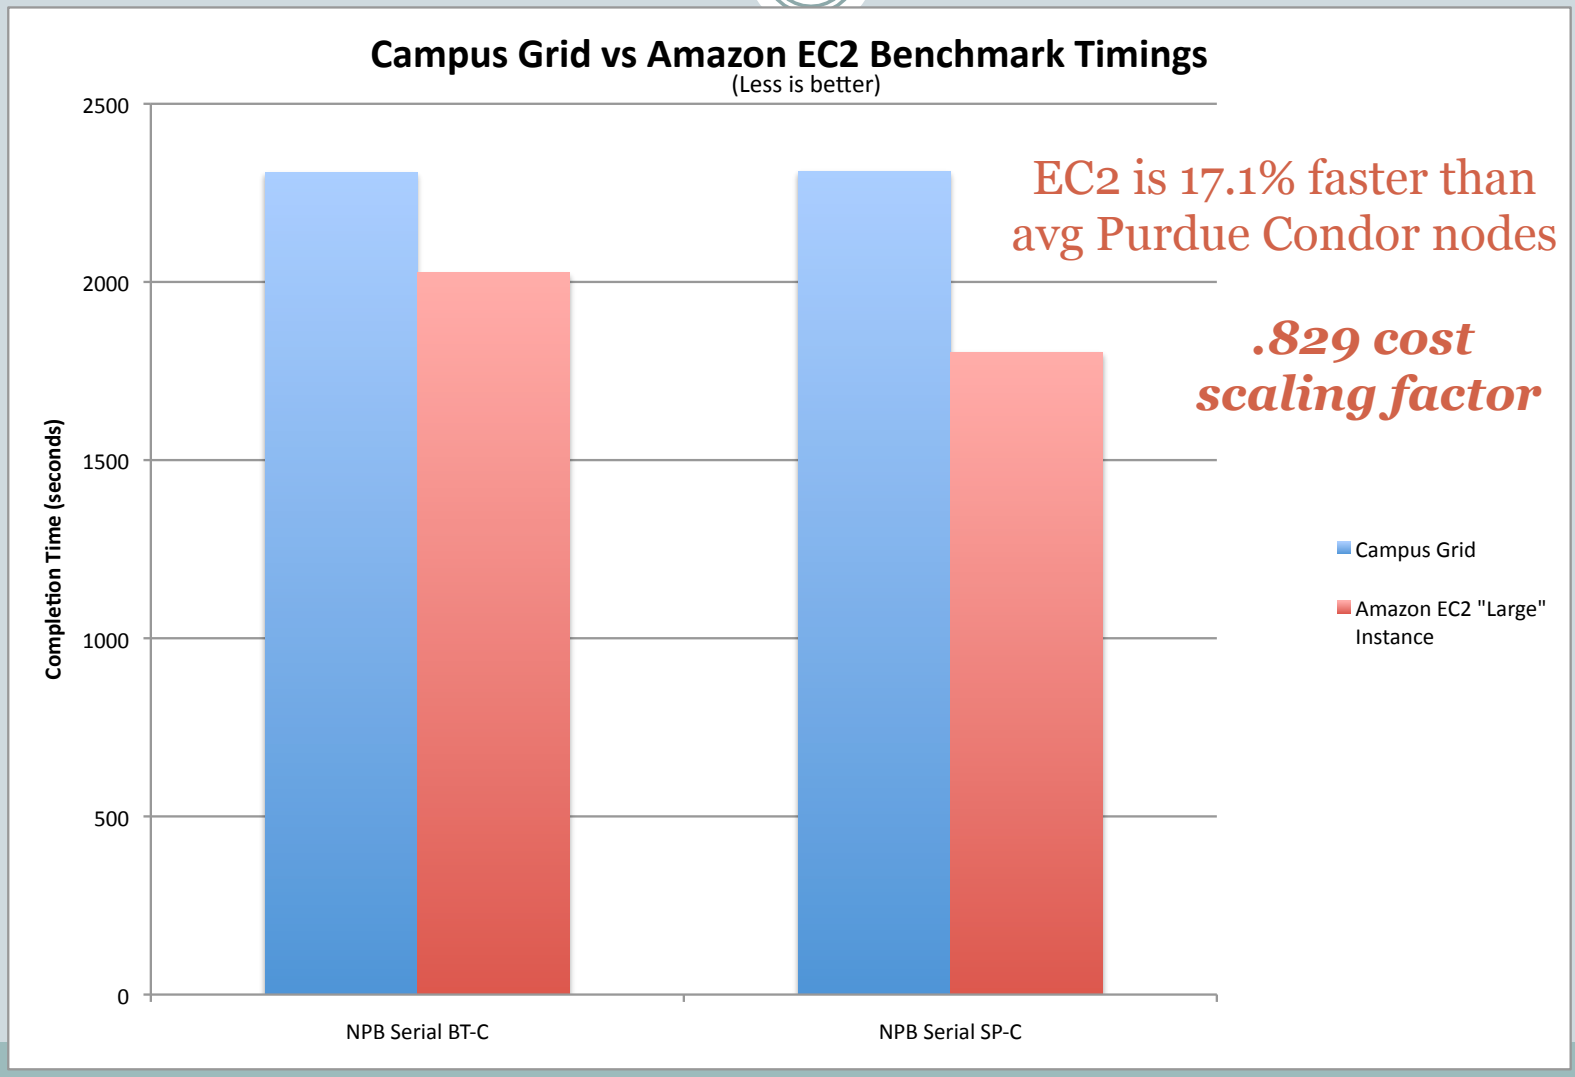

Normalized Cost Model

$$C_n = \frac{C_h}{F_n},$$

$$\begin{aligned} C_n &= \frac{C_h}{F_n} \\ &= \frac{.0300}{.829} \\ &= \$0.03619, \end{aligned}$$

C_n: normalized per-core-hour cost
C_h: pre-normalized per core-hour cost **$.300**
F_n: a constant representing the normalizing factor of one hour on the grid to 1 hour on EC2. **.829**

Normalized core-hour cost: $.03619

Additional Costs

Item	Total yearly Cost
Systems Engineering (1 FTE)	$73,810.00
User Support (.75 FTE)	$55,357.50
Distributed IT Staff (.1 FTE)	$11,071.50
Additional Power Load	$290,295.01
	Amortized Over 5 Years
Submit Nodes	$6,360.00
Checkpoint Servers ,etc	$8,480.00
Total	$433,502.01

Additional Costs – Per Core Hour

$$C_a = \frac{E_y * S_a}{H_y},$$

$$\begin{aligned} C_a &= \frac{E_y * S_a}{H_y} \\ &= \frac{\$433,502.01 * 13,526}{8760} \\ &= \$0.003658623 \text{ (3.66 tenths of a cent).} \end{aligned}$$

C_a: additional per-core-hour cost
E_y: total yearly additional cost of operating the grid — $433,502.01
S_a: total available slots in the grid — 13,526
H_y: total hours in a year — 8760

3.66 tenths of one cent!

Power Costs

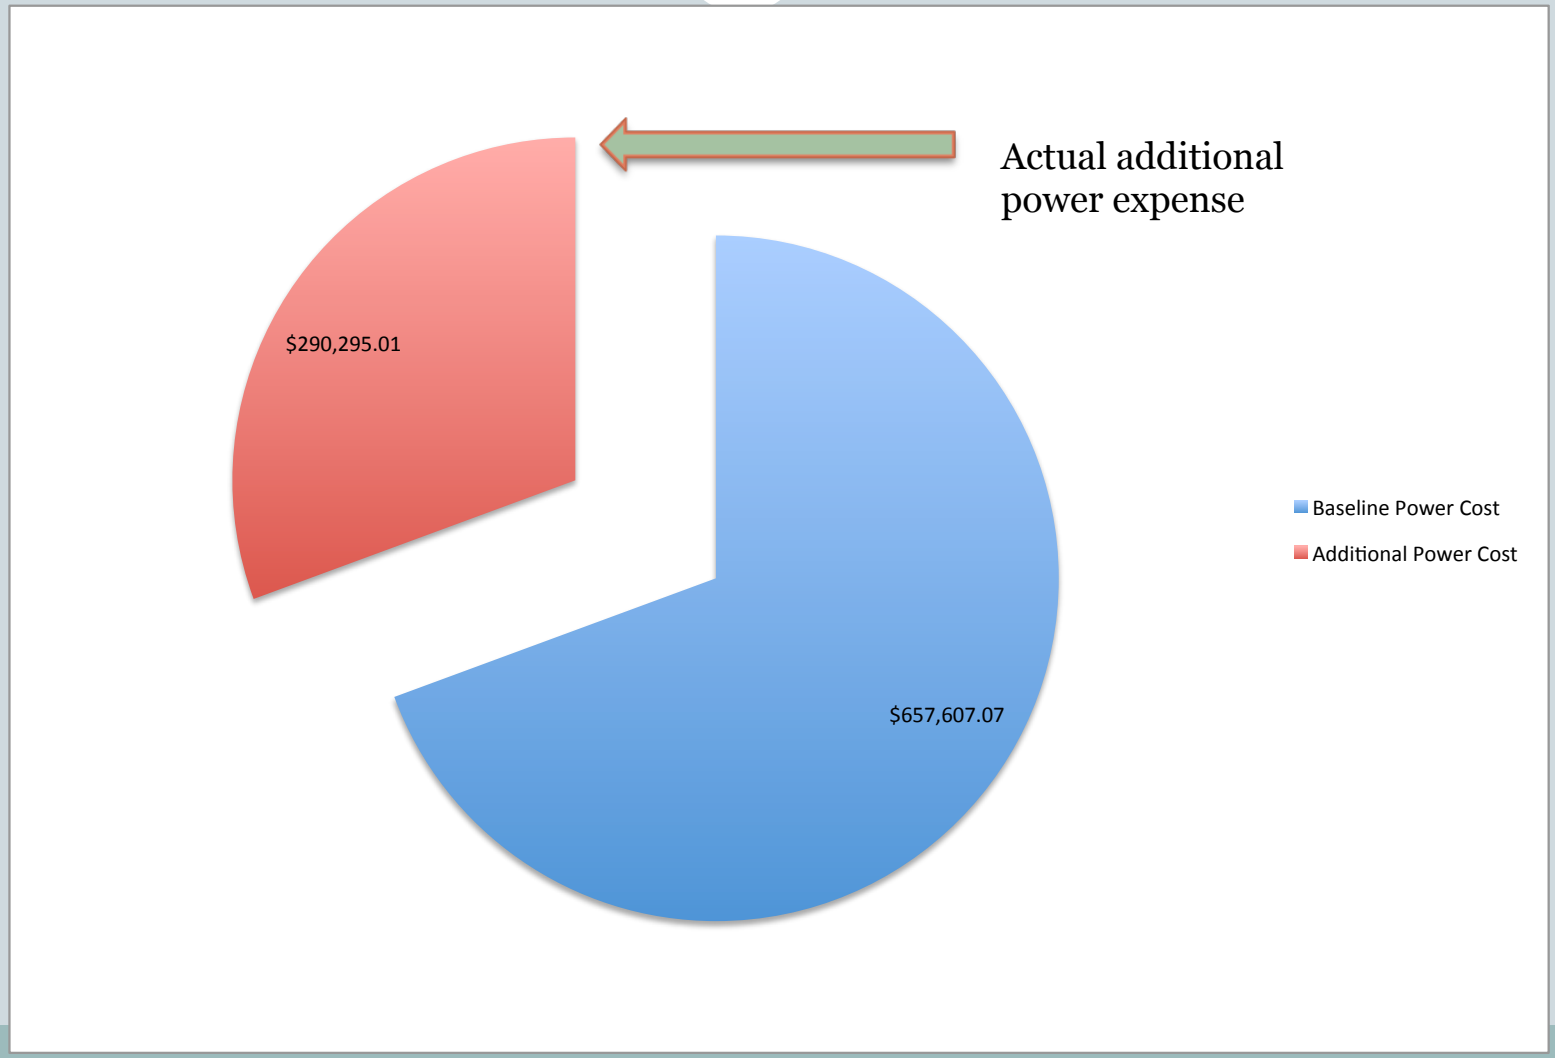
Actual additional power expense
$290,295.01
$657,607.07
Baseline Power Cost
Additional Power Cost

Total Cost

$$C_t = C_n + C_a$$

$$= \$0.03619 + \$0.003658623$$

$$= \$0.0398468012$$

C_t: total per-core-hour cost
C_n: total normalized base cost of the campus grid
C_a: total additional cost per core hour

Total per core-hour cost: $.03985

Scientific Output – Raw Metrics

Year	Unique users	Unique Pis	Unique PI Depts	Fields of Science	Jobs	Hours
2005	25	8	5	4	295265	1.9 M
2006	70	27	11	11	4.44 M	4.61 M
2007	115	50	16	19	9.93 M	8.17 M
2008	115	60	13	18	14.9 M	16.6 M
2009	163	85	18	16	15.4 M	17.9 M
2010	145	79	20	16	15.2 M	18.6 M

From Rosen Center Usage Metrics

Solutions or Publications as Metrics

- Solutions per unit of time is one metric recommended in the literature
  - How much ***good*** computation was done in those millions of hours?
  - But, from the perspective of the institution, this is hard to obtain
    - Only the user knows how many of these jobs were scientifically useful!
- Publications are the end goal of research, so they are an excellent measure of output
  - Unfortunately no data exists on publications directly attributable to the campus grid

Per-Metric Costs

$$C_{tot} = \frac{H_y * C_t}{M_u},$$

2005:

$$\begin{aligned} C_{tot} &= \frac{H_y * C_t}{M_u} \\ &= \frac{(1,945,723 \text{ hours} * \$0.0398468012)}{25 \text{ Unique Users}} \\ &= \$3,101.23 \text{ per user} \end{aligned}$$

C_{tot}: the total per-core-hour cost per unit of M_u
H_y: total hours provided in a year
C_t: total cost of an hour of use in the campus grid
M_u: metric of use (such as users or PIs)

1,945,723
$0.0398468012
users

$3, 101.23 per user

Additional Costs, Per Metric

	Per User	Per PI	Per PI Dept	Per Field of Science
2005	$284.75	$889.83	$1,423.73	$1,779.67
2006	$241.13	$625.14	$1,534.45	$1,534.45
2007	$259.87	$597.69	$1,867.78	$1,572.87
2008	$528.17	$1,312.33	$4,662.24	$3,374.40
2009	$403.96	$774.66	$3,658.12	$4,115.39
2010	$469.54	$861.81	$3,404.15	$4,255.18
	Per User	**Per PI**	**Per PI Dept**	**Per Field of Science**
Average	$364.57	$793.58	$2,760.08	$2,771.99

$364.57 a user

Average 11.3 Million Hours

Average 105 users

Average 107,300 hours per user

Additional Costs, Per Metric

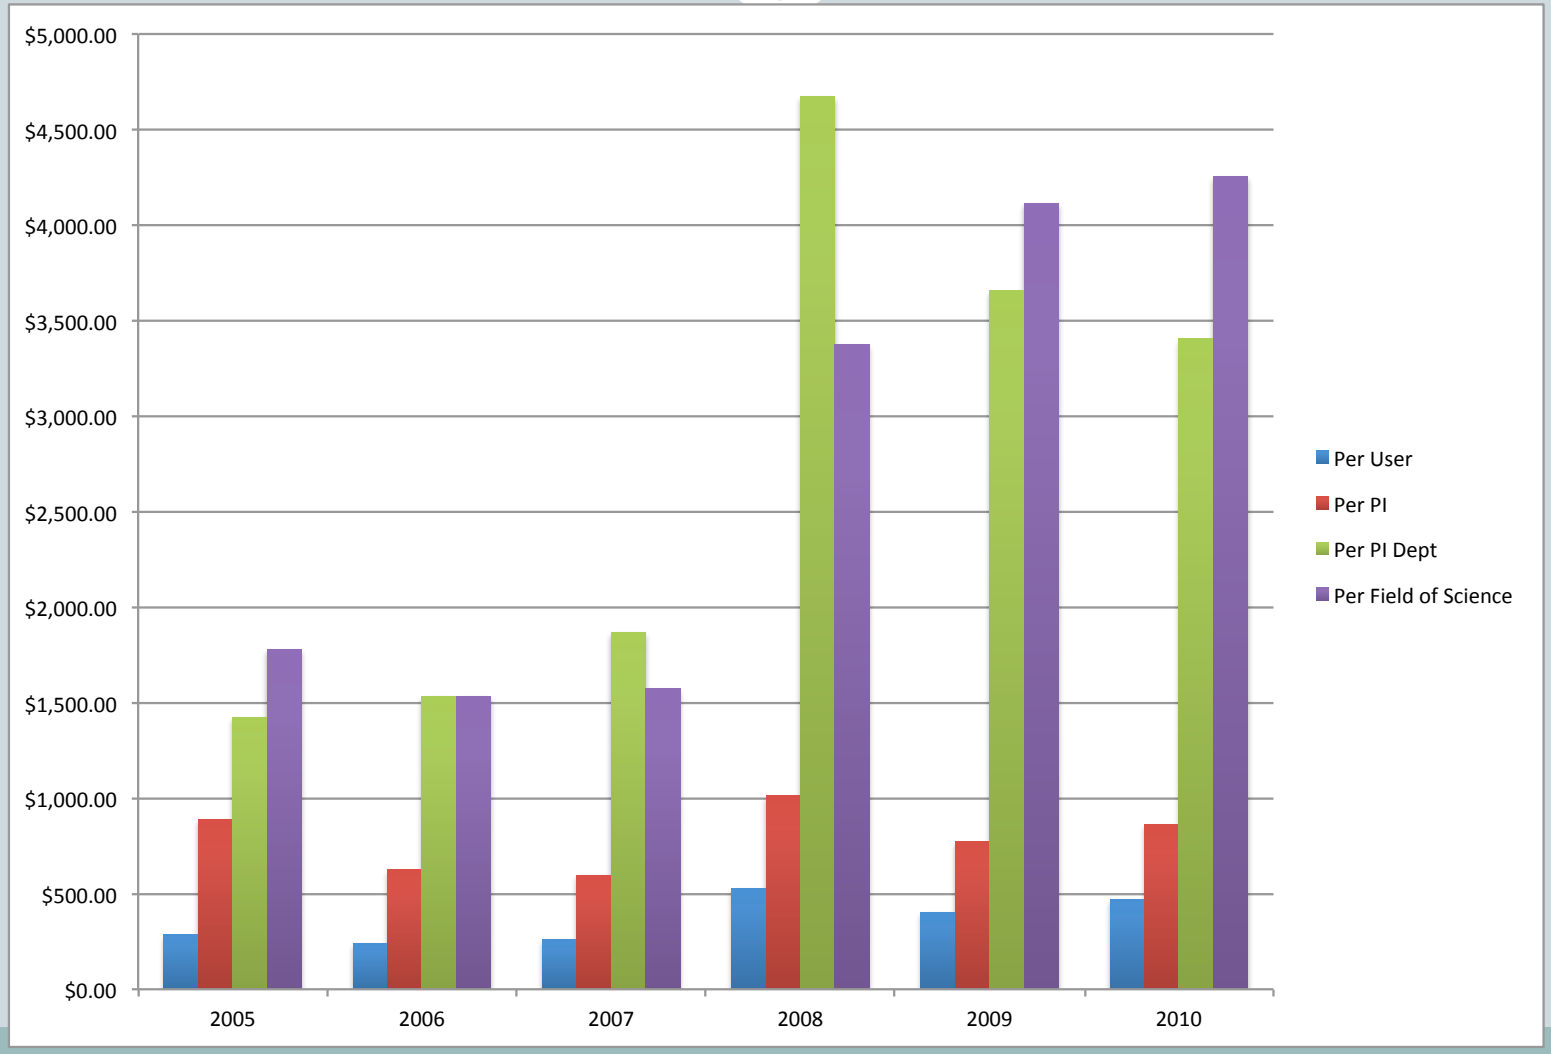

Summary

- Measured Relative Performance of Grid Nodes
  - .829 relative to Amazon EC2
- Developed Models and Calculated Per-Core Hour Costs:
  - Normalized: $.03619
  - Additional: $0.003658623
  - Total: $0.039847
- Calculated Costs per unit of Several Metrics
  - For example: For each user of the grid in 2010
    - Additional cost to Purdue is $469.54
- On average, each user costs Purdue an extra $364.57

Recommendation

- A campus grid is indeed a cost-effective way to create a useful HPC resource
  - Any institution with a substantial investment in an IT infrastructure should consider a campus grid to support HPC
- Questions?

The End

- Questions?